

LVPNet: A Latent-variable-based Prediction-driven End-to-end Framework for Lossless Compression of Medical Images

Chenyue Song¹, Chen Hui^{1,2,*}, Qing Lin³, Wei Zhang¹, Siqiao Li¹, Haiqi Zhu¹, Shengping Zhang¹, Zhixuan Li⁴, Shaohui Liu¹, Feng Jiang², Xiang Li¹

¹Harbin Institute of Technology, China

²Nanjing University of Information Science and Technology, China

³Dalian University of Technology, China

⁴Nanyang Technological University, Singapore
cysong@stu.hit.edu.cn, chui@nuist.edu.cn

Abstract. Autoregressive Initial Bits is a framework that integrates sub-image autoregression and latent variable modeling, demonstrating its advantages in lossless medical image compression. However, in existing methods, the image segmentation process leads to an even distribution of latent variable information across each sub-image, which in turn causes posterior collapse and inefficient utilization of latent variables. To deal with these issues, we propose a prediction-based end-to-end lossless medical image compression method named LVPNet, leveraging global latent variables to predict pixel values and encoding predicted probabilities for lossless compression. Specifically, we introduce the Global Multi-scale Sensing Module (GMSM), which extracts compact and informative latent representations from the entire image, effectively capturing spatial dependencies within the latent space. Furthermore, to mitigate the information loss introduced during quantization, we propose the Quantization Compensation Module (QCM), which learns the distribution of quantization errors and refines the quantized features to compensate for quantization loss. Extensive experiments on challenging benchmarks demonstrate that our method achieves superior compression efficiency compared to state-of-the-art lossless image compression approaches, while maintaining competitive inference speed. The code is at <https://github.com/scy-Jackel/LVPNet>.

Keywords: Lossless medical image compression · Latent variable · Pixel prediction · Quantization.

* Chen Hui is corresponding author.

This work was supported in part by the Central Guidance for Local Science and Technology Development Fund Projects under grant 2024ZYD0266, in part by the Startup Foundation for Introducing Talent of Nanjing University of Information Science and Technology under grant 2025r029.

1 Introduction

In medical imaging applications, the handling, transmission, and backup of data often necessitate the storage of lossless images, resulting in substantial storage demands [1–3]. Lossless image compression techniques play a crucial role in mitigating storage costs by leveraging inherent image correlations to reduce file sizes while maintaining the integrity of the original content. Shannon’s source coding theorem [4] defines the lower bound for code length based on image entropy, which is determined by the image’s inherent distribution. Thus, reducing redundant information is a fundamental strategy for image compression. [5] introduce the use of the LZ77 algorithm [6] and Huffman coding [7] to eliminate redundant information. Later, reference [8] propose the discrete wavelet transform as an effective method for capturing multi-scale image features.

With the advancement of deep learning, various models [9–11] have been widely adopted for lossless image compression. Representative methods include autoregressive models (ARMs) [12–14], variational autoencoders (VAEs) [15–20], normalizing flows (NFs) [21–24], prediction model [25–28], as well as diffusion models [29, 30]. These methods require balancing compression performance with inference speed. For example, ARMs and diffusion models offer impressive compression performance, but come with the trade-off of significantly longer inference times [31–33]. Moreover, image compression encompasses both dataset-level and single-image compression. However, existing lossless compression methods are typically tailored for specific data types, facing significant challenges in addressing both application scenarios simultaneously. For instance, while NFs excel in dataset compression, they exhibit relatively weaker performance in single-image compression. Additionally, as the size of the autoregressive prior increases in ARMs, the risk of posterior collapse along the sub-image sequence becomes more pronounced [34–36], which negatively affects compression performance.

To deal with these challenges, we propose LVPNet, a latent-variable-based prediction-driven end-to-end lossless compression framework specially designed for medical images, suitable for both datasets and single images. By extracting effective low-dimensional latent representations from the entire image, our approach efficiently captures spatial dependencies within the latent space and predicts pixel values in this domain. Specially, we introduce two key components: First, Global Multi-scale Sensing Module (GMSM) captures a variety of multi-scale features during downsampling and encodes relevant information into the latent variables; Then Quantization Compensation Module (QCM) learns the distribution of quantization errors to compensate for the loss and enhance pixel prediction accuracy. Our contributions are summarized as follows:

- We propose latent-variable-based prediction-driven end-to-end framework for lossless compression of medical images, applicable to both dataset and single-image compression tasks.
- We propose the GMSM and QCM modules to extract more comprehensive information from multi-level features, which helps mitigate the issue of posterior collapse and reduces the information loss from quantization.

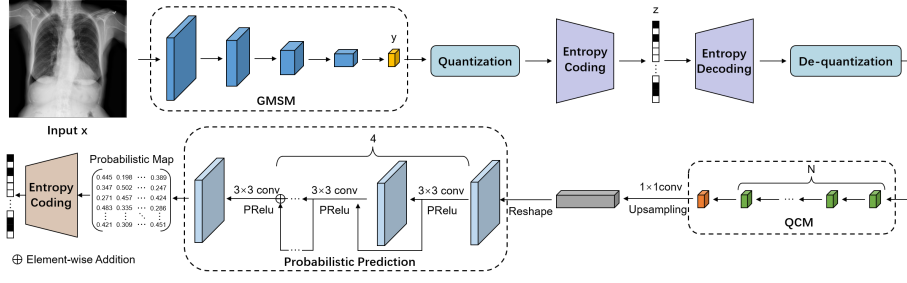


Fig. 1: Overview of the proposed framework. The framework consists of five distinct modules: a global multi-level Sensing module, quantization, entropy coding, a quantization compensation module, and a probability prediction module.




Fig. 2: The structure of GSM and QCM. (a) r is the pre-defined sampling rate. (b) We employ N residual connections and a global connection to propagate features from various layers to the final output.

- The proposed model exhibits exceptional inference speed and achieves strong compression performance across benchmark datasets for medical imaging.

2 Methodology

2.1 Overview of LVPNet

The proposed framework is shown in Fig. 1. The input image x is first processed by the global multi-level perception module G to generate the compressed coefficients $y = G(x)$, with a sampling ratio of r . y is quantized to map it to the integer domain, producing the latent variable $z = Q(y)$, which is then entropy encoded into a bitstream for storage. Afterward, z undergoes dequantization to produce $\hat{y} = Deg(z)$, and the quantization loss is compensated by the QCM. The compensated features are upsampled and reconstructed, after which the pixel values for the entire image are predicted using the probability prediction module. The process from z to predicted probabilities is denoted as $p_\theta(x|z)$, where θ represents the model parameters. Finally, the predicted probabilities are entropy encoded into a bitstream for storage. LVPNet only needs to store the entropy-encoded z , the bitstream of predicted probabilities $p_\theta(x|z)$, and the model parameters θ . Since θ has a fixed size and the bitstream of z remains relatively stable due to

the fixed sampling ratio r , LVPNet focuses on reducing the size of the entropy-encoded bitstream for $p_\theta(x|z)$. During inference, the latent variable z is decoded, and the predicted pixel values are generated using $p_\theta(x|z)$. By comparing these predictions with the entropy-decoded probabilities corresponding to the ground truth pixel values, the image can be losslessly reconstructed.

2.2 Global Multi-scale Sensing Module

Convolutional neural networks hierarchically extract features, where layers proximal to the input capture low-level features such as edges and basic textures, while deeper layers encode abstract and semantic features. Guided by these principles, we propose the GSM method, as illustrated in Figure 1. The framework operates in two stages. In the first stage, 2×2 convolutional layers are employed to extract features. In the second stage, features across all hierarchical levels are aggregated and sampled using 1×1 convolutional layers, instead of exclusively relying on low-level or high-level features for sampling.

In GSM, to aggregate multi-level features for sampling, we draw inspiration from ResNet [37], employing skip connections to propagate features from various layers to the final stage. Additionally, pooling layers are incorporated to ensure dimensionality alignment. The corresponding formula is expressed as follows:

$$f_{t+1} = \text{Conv}(f_t) + \text{Pool}(f_t) \quad (1)$$

where f_t and f_{t+1} denote the sampling features of the t -th and $(t+1)$ -th layers, respectively. When transforming the current feature $f_t \in \mathbb{R}^{C \times H \times W}$ into the next-layer feature $f_{t+1} \in \mathbb{R}^{4C \times \frac{H}{2} \times \frac{W}{2}}$, the feature volume remains consistent. Downsampling is performed through the 1×1 convolution layer in the second stage, which facilitates adaptation to various sampling ratios.

2.3 Quantization and Entropy Coding

Quantization The downsampled values y are quantized to produce the discrete latent variables z , and the inverse operation is applied during dequantization to recover the reconstructed coefficient \hat{y} :

$$z = Q(y) = \lfloor y/Q_{step} \rfloor, \hat{y} = Deq(z) = z \times Q_{step} \quad (2)$$

where Q_{step} denotes the quantization step size. Due to quantization causing the gradients to become zero, preventing effective backpropagation to the GSM, we store the gradient prior to the floor operation during forward propagation and use this stored gradient for backpropagation updates. The gradient update formula for the loss function \mathcal{L} in the quantization module is as follows:

$$\frac{\partial \mathcal{L}}{\partial y} = \frac{\partial \mathcal{L}}{\partial z} \cdot \frac{\partial z}{\partial y} \approx \frac{\partial \mathcal{L}}{\partial z} \cdot \frac{1}{Q_{step}} \quad (3)$$




Fig. 3: The feature map visualization of the GSM and QCM modules. (a) Original image; (b) f_1 in Eq. 1; (c) f_2 in Eq. 1; (d) f_3 in Eq. 1; (e) Output of GSM y ; (f) Output of QCM; (g) Residual between y and QCM output.

Entropy Coding We apply entropy encoding to both the latent variable z and the predicted probabilities of the pixel values in $p_\theta(\hat{x}|z)$. A well-trained model typically results in highly sparse residual latent variables. To encode these sparse residuals, we use Huffman encoding [38], while the predicted probabilities are encoded using arithmetic coding [39].

2.4 Quantization Compensation Module

The quantization compensation module helps mitigate quantization loss and enhance the predicted probabilities. We employ a basic CNN architecture for the compensation module, as illustrated in Figure 1. This network incorporates a global shortcut and multiple residual blocks, which accelerate training. During quantization, a floor operation is applied, causing the dequantized measurement \hat{y} to be slightly smaller than the original measurement y . To address this, we use the $\text{relu}(\cdot)$ activation function in the final layer of the compensation module, ensuring that the compensation value remains non-negative.

2.5 Loss Function

We define the loss function for lossless compression of datasets as follows:

$$\mathcal{L}(\theta, G, Q) = -\mathbb{E}_{x \sim \mathcal{P}_{\text{data}}} [\log_2 (p_\theta(\hat{x} | Q(G(x))))] \quad (4)$$

$$= -\sum_{k=1}^N \sum_{i=1}^{H^{(k)}} \sum_{j=1}^{W^{(k)}} \delta(x_{ij}^{(k)}, 1) \cdot \left(\log_2 p_\theta(\hat{x}_{ij}^{(k)} | Q(G(x^{(k)}))) \right) \quad (5)$$

where N denote the number of images in the dataset, with $W^{(k)}$ and $H^{(k)}$ representing the width and height of the k -th image. The true pixel value at the

Table 1: Comparison of dataset compression performance in BPP. Group 1: Traditional methods, Group 2: Methods based on machine learning.

Method	Chest X-ray	CIFAR10	ImageNet32	ImageNet64
PNG [5]	3.07	5.89	6.39	5.74
FLIF [40]	2.87	4.19	4.52	4.54
JPEG-XL [41]	2.93	5.74	6.39	5.89
L3C [31]	2.89	-	4.76	4.42
IDF [22]	2.78	3.34	4.18	3.90
Hilloc [18]	3.02	4.56	4.20	3.90
SHVC [32]	2.83	3.16	3.98	3.68
iVPF [24]	2.80	3.20	4.03	3.75
LBB [21]	2.76	3.12	3.88	3.70
iFlow [23]	2.77	3.12	3.87	3.70
BCM-Net [28]	2.87	3.26	4.06	4.19
HMEM [20]	2.79	3.19	3.89	3.80
ArIB-BPS [14]	2.74	3.06	3.91	3.63
LVPNet (ours)	2.65±0.03	3.02±0.02	3.76±0.04	3.51±0.03

i -th row and j -th column of the k -th image is one-hot encoded as $x_{ij}^{(k)}$. $\mathcal{P}_{\text{data}}$ denotes the true data distribution, and δ is the Kronecker delta function. The loss function quantifies the storage cost associated with storing the predicted probabilities, and for lossless compression of a single image it's defined as:

$$\mathcal{L}(\theta, G, Q) = - \sum_{i=1}^H \sum_{j=1}^W \delta(x_{ij}, 1) \cdot (\log_2 p_{\theta}(\hat{x}_{ij} \mid Q(G(x)))) \quad (6)$$

3 Experiments

3.1 Training Details

We evaluate the performance of our method for both dataset-level and single-image lossless compression on Chest X-ray[42], CIFAR10, ImageNet32, and ImageNet64. Of these, Chest X-ray is a medical image dataset, while the others are natural image datasets. We use bits per pixel (BPP) as the evaluation metric, where lower values indicate better performance.

We use the Adam optimizer [43] with an initial learning rate of 0.0001 and step decay for scheduling. The GSM sampling rate r is set to 0.15, and the quantization step size Q_{step} is 0.01. For dataset compression, the QCM consists of 24 residual blocks, while for single image compression, it contains only 3 residual blocks. All experiments are conducted on an NVIDIA GeForce RTX 3060.

3.2 Visualization of Feature Maps

In this section, we visualize the feature maps of the GSM and QCM modules, as shown in Figure 3. From Figures 3b to 3e, we observe that, during GSM's

Table 2: Comparison of single image compression performance in BPP. Δ : These values are derived by summing the dataset’s compression performance and the initial bit count.

Method	Chest X-ray	CIFAR10	ImageNet32	ImageNet64
PNG [5]	3.07	5.89	6.39	5.74
FLIF [40]	2.87	4.19	4.52	4.54
JPEG-XL [41]	2.93	5.74	6.39	5.89
L3C [31]	2.89	-	4.76	4.42
IDF [22]	2.78	3.34	4.18	3.90
iVPF [24]	5.83 Δ	9.20 Δ	-	-
LBB [21]	37.53 Δ	42.98 Δ	49.84 Δ	41.70 Δ
iFlow [23]	29.89 Δ	37.40 Δ	38.27 Δ	38.12 Δ
BCM-Net [28]	3.21	4.03	4.10	4.18
HMEM [20]	3.15	3.97	4.04	4.13
ArIB-BPS [14]	2.76	3.07	3.92	3.64
LVPNet (ours)	2.72\pm0.02	3.10 \pm 0.02	3.82\pm0.03	3.57\pm0.03

Table 3: Comparison of inference times.

Dataset	Method	Compression Performance(BPP) \downarrow		Inference Time(ms/sample) \downarrow	
		dataset	single	encode	decode
Chest X-ray	LBB [21]	2.762	37.529 Δ	643.95	634.19
	iFlow [23]	2.768	29.894 Δ	451.87	512.48
	ArIB-BPS [14]	2.741	2.763	197.24	203.50
	LVPNet (ours)	2.650	2.719	143.63	184.06
CIFAR10	LBB [21]	3.118	49.835	64.94	64.94
	iFlow [23]	3.118	37.398	19.38	47.56
	ArIB-BPS [14]	3.057	3.070	14.93	14.93
	LVPNet (ours)	3.024	3.099	16.44	20.28
ImageNet32	LBB [21]	3.875	49.835	194.14	194.14
	iFlow [23]	3.873	38.273	74.84	119.30
	ArIB-BPS [14]	3.911	3.918	52.96	52.96
	LVPNet (ours)	3.758	3.822	39.73	48.17


downsampling process, the earlier layers effectively capture the skeletal and organ contours in the X-ray image, while the later layers extract more abstract features. This observation suggests that GSM progressively extracts key information in a coarse-to-fine manner. As shown in Figure 3f, the output from the QCM module closely resembles the feature map of y in Figure 3e. Additionally, the residual map in Figure 3g appears generally darker, indicating the QCM module’s effectiveness in compensating for quantization loss.

3.3 Compression Performance


We compare our method with traditional approaches and state-of-the-art deep learning methods. We present the detailed results of the experiment in Table 1. Our method achieves superior compression performance, particularly in compressing medical image datasets, while also maintaining competitive performance

Table 4: Ablation study of the effectiveness of GSM and QCM.

Sampling	QCM	Compression Performance(BPP)↓		Inference Time(ms/sample)↓	
		dataset	single	encode	decode
CNN		2.833	2.906	133.11	166.89
CNN	✓	2.757	2.859	132.15	184.12
GMSM		2.732	2.841	143.52	167.03
GMSM	✓	2.650	2.719	143.63	184.06



(a) Dataset compression



(b) Single image compression

Fig. 4: Ablation study for different sampling rates r .

on natural image datasets. ArIB-BPS shows limited compression capability due to issues like posterior collapse and inefficient use of latent variables, highlighting GSM’s effectiveness in leveraging global latent variables for prediction.

To further validate the effectiveness of our method, we evaluate its performance on the lossless compression of individual images, with results presented in Table 2. The experimental findings indicate that our method achieves high compression efficiency in single-image compression tasks. Compared to state-of-the-art lossless compression methods, it consistently outperforms existing approaches across multiple datasets. These results further demonstrate the applicability and robustness of our method, proving its suitability for both large-scale dataset compression and single-image lossless compression.

In addition, we perform a comprehensive comparison between our method and state-of-the-art approaches, which demonstrate strong performance in both large-scale image datasets and single-image compression tasks. As shown in Table 3, LVPNet achieves outstanding performance across various datasets while achieving an inference speed on par with the fastest existing approach, indicating a better balance between compression efficiency and inference speed.

3.4 Ablation Studies

We perform three ablation experiments on the Chest X-ray dataset to demonstrate the efficacy of our latent-variable-based, prediction-driven model for medical image lossless compression and the effectiveness of discrete sampling. We assess the impact of autoregression across different dimensions by training and evaluating three additional models: one in which the global multi-level perception module is replaced by a Convolutional Neural Network (CNN), one in which

the quantization compensation network is removed, and one where both modifications are applied. As shown in the experimental results in Table 4, our model achieves superior compression performance while maintaining comparable inference time. Additionally, we evaluate the effectiveness of our customized sampling rate strategy. As shown in Figure 4, this sampling rate provides better compression performance, striking a favorable balance between latent variable storage overhead and compression prediction accuracy.

4 Conclusion

We propose a novel method for lossless medical image compression, LVPNet. The proposed GSM module allows us to make full use of latent variables, effectively modeling the spatial dependencies within the latent space. In addition, we introduce a quantization compensation module, which corrects quantized features by modeling the distribution of quantization errors, thereby mitigating the effects of quantization loss. Experimental results demonstrate that, compared to the current state-of-the-art methods in lossless image compression, our method achieves superior compression performance with comparable inference time.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Li, Junjian, et al. "CA 2 CL: Cluster-Aware Adversarial Contrastive Learning for Pathological Image Analysis." *IEEE Journal of Biomedical and Health Informatics* (2025).
2. Xu, Jiawei, et al. "Semantic-Orthogonal Multi-modal Attention Network for RGB-D Salient Object Detection." *The Visual Computer* (2025): 1-13.
3. Li, Junjian, et al. "DARC: Deep adaptive regularized clustering for histopathological image classification." *Medical image analysis* 80 (2022): 102521.
4. Shannon, Claude Elwood. "A mathematical theory of communication." *The Bell system technical journal* 27.3 (1948): 379-423.
5. Boutell, Thomas. Png (portable network graphics) specification version 1.0. No. rfc2083. 1997.
6. Ziv, Jacob, and Abraham Lempel. "A universal algorithm for sequential data compression." *IEEE Transactions on information theory* 23.3 (1977): 337-343.
7. Huffman, David A. "A method for the construction of minimum-redundancy codes." *Proceedings of the IRE* 40.9 (1952): 1098-1101.
8. "JPEG 2000 standard ISO/IEC 15444 ITU-T Recommendation T.800", Joint Photographic Experts Group (JPEG), 2004, [online] Available: <http://www.jpeg.org/jpeg>.
9. Bond-Taylor, Sam, et al. "Deep generative modelling: A comparative review of vaes, gans, normalizing flows, energy-based and autoregressive models." *IEEE transactions on pattern analysis and machine intelligence* 44.11 (2021): 7327-7347.

10. Yu, Xinlei, et al. "ICH-PRNet: a cross-modal intracerebral haemorrhage prognostic prediction method using joint-attention interaction mechanism." *Neural Networks* 184 (2025): 107096.
11. Luo, Haozhe, et al. "On the Interplay of Human-AI Alignment, Fairness, and Performance Trade-offs in Medical Imaging." *arXiv preprint arXiv:2505.10231* (2025).
12. Salimans, Tim, et al. "Pixelcnn++: Improving the pixelcnn with discretized logistic mixture likelihood and other modifications." *arXiv preprint arXiv:1701.05517* (2017).
13. Van Den Oord, Aaron, Nal Kalchbrenner, and Koray Kavukcuoglu. "Pixel recurrent neural networks." *International conference on machine learning*. PMLR, 2016.
14. Zhang, Zhe, et al. "Learned Lossless Image Compression based on Bit Plane Slicing." *CVPR*. 2024.
15. Kingma, Friso, Pieter Abbeel, and Jonathan Ho. "Bit-swap: Recursive bits-back coding for lossless compression with hierarchical latent variables." *International Conference on Machine Learning*. PMLR, 2019.
16. Mentzer, Fabian, et al. "Practical full resolution learned lossless image compression." *CVPR*. 2019.
17. Townsend, James, Tom Bird, and David Barber. "Practical lossless compression with latent variables using bits back coding." *arXiv preprint arXiv:1901.04866* (2019).
18. Townsend, James, et al. "Hilloc: Lossless image compression with hierarchical latent variable models." *arXiv preprint arXiv:1912.09953* (2019).
19. Bai, Yuanchao, et al. "Deep lossy plus residual coding for lossless and near-lossless image compression." *TPAMI* 46.5 (2024): 3577-3594.
20. Fu, Chuan, Bo Du, and Liangpei Zhang. "Hybrid-context-based multi-prior entropy modeling for learned lossless image compression." *Pattern Recognition* 155 (2024): 110632.
21. Ho, Jonathan, Evan Lohn, and Pieter Abbeel. "Compression with flows via local bits-back coding." *Advances in Neural Information Processing Systems* 32 (2019).
22. Hoogeboom, Emiel, et al. "Integer discrete flows and lossless compression." *Advances in Neural Information Processing Systems* 32 (2019).
23. Zhang, Shifeng, et al. "iflow: Numerically invertible flows for efficient lossless compression via a uniform coder." *Advances in Neural Information Processing Systems* 34 (2021): 5822-5833.
24. Zhang, Shifeng, et al. "ivpf: Numerical invertible volume preserving flow for efficient lossless compression." *CVPR*. 2021.
25. Ulacha, Grzegorz, and Mirosław Łazoryszczak. "Lossless Image Compression Using Context-Dependent Linear Prediction Based on Mean Absolute Error Minimization." *Entropy* 26.12 (2024): 1115.
26. Rhee, Hochang, et al. "Lossless image compression by joint prediction of pixel and context using duplex neural networks." *IEEE Access* 9 (2021): 86632-86645.
27. Rahman, Md Atiqur, and Mohamed Hamada. "A prediction-based lossless image compression procedure using dimension reduction and Huffman coding." *Multimedia Tools and Applications* 82.3 (2023): 4081-4105.
28. Liu, Xiangrui, et al. "Bilateral context modeling for residual coding in lossless 3D medical image compression." *TIP* (2024).
29. Hoogeboom, Emiel, et al. "Autoregressive diffusion models." *arXiv preprint arXiv:2110.02037* (2021).
30. Kingma, Diederik, et al. "Variational diffusion models." *Advances in neural information processing systems* 34 (2021): 21696-21707.

31. Mentzer, Fabian, et al. "Practical full resolution learned lossless image compression." CVPR. 2019.
32. Ryder, Tom, et al. "Split hierarchical variational compression." CVPR. 2022.
33. Yang, Yanwu, et al. "Creg-kd: Model refinement via confidence regularized knowledge distillation for brain imaging." Medical Image Analysis 89 (2023): 102916.
34. Lucas, James et al. "Don't Blame the ELBO! A Linear VAE Perspective on Posterior Collapse." ArXiv abs/1911.02469 (2019): n. pag.
35. Luo, Haozhe, Yu Changdong, and Raghavendra Selvan. "Hybrid ladder transformers with efficient parallel-cross attention for medical image segmentation." International conference on medical imaging with deep learning. PMLR, 2022.
36. Yang, Yanwu, et al. "Advancing Brain Imaging Analysis Step-by-Step via Progressive Self-paced Learning." International Conference on Medical Image Computing and Computer-Assisted Intervention. Cham: Springer Nature Switzerland, 2024.
37. He, Kaiming, et al. "Deep residual learning for image recognition." CVPR. 2016.
38. Huffman, David A. "A method for the construction of minimum-redundancy codes." Proceedings of the IRE 40.9 (1952): 1098-1101.
39. B. P. Tunstall, "Synthesis of noiseless compression codes," in PhD diss., Georgia Institute of Technology, 1967.
40. Sneyers, Jon, and Pieter Wuille. "FLIF: Free lossless image format based on MANNIAC compression." ICIP. IEEE, 2016.
41. Alakuijala, Jyrki, et al. "JPEG XL next-generation image compression architecture and coding tools." Applications of digital image processing XLII. Vol. 11137. SPIE, 2019.
42. Wang, Xiaosong, et al. "Chestx-ray8: Hospital-scale chest x-ray database and benchmarks on weakly-supervised classification and localization of common thorax diseases." CVPR. 2017.
43. Kingma, Diederik P. and Jimmy Ba. "Adam: A Method for Stochastic Optimization." CoRR abs/1412.6980 (2014): n. pag.